

Aggregate Implications of Capacity Constraints under Demand Shocks

Chacko George and Florian Kuhn *

University of Texas at Austin

April 29, 2014

Abstract

We investigate the impact of firm capacity constraints on aggregate production and productivity when the economy is driven by aggregate and idiosyncratic demand shocks. We are motivated by three observed regularities in US data: business cycles are asymmetric, in that large absolute changes in output are more likely to be negative than positive; capacity and capital utilization are procyclical, and increase the procyclicality of measured productivity; the dispersion of firm productivity increases in recessions.

We devise a model of demand shocks and endogenous capacity constraints that is qualitatively consistent with these observations. We then calibrate the model to aggregate utilization data using standard Bayesian techniques. Quantitatively, we find that the calibrated model also exhibits significant asymmetry in output, on the order of the regularities observed in GDP.

*Department of Economics, University of Texas at Austin, 2225 Speedway, C3100, Austin TX, 78712, email: chacko.george@utexas.edu, florian.kuhn@utexas.edu.

1 Introduction

Despite their (surprising) symmetry, two very asymmetric patterns have been observed in business cycles. First, big booms occur less often than big busts, that is, large absolute changes in output are more likely to be negative than positive. Second, firm productivity becomes more dispersed in recessions, and the productivity of the worst firms decreases relative to the mean (Kehrig (2011)).

We believe these patterns are closely related to the way firms change the intensity of their production processes over the business cycle. Capacity utilization, or the percentage of potential output that firms produce, seems to be procyclical, as does capital utilization, or the intensity with which firms use their machinery. It has also been noted that, even though average productivity is apparently procyclical, this is no longer the case when we take the changes in capacity utilization into account (Basu et al. (2006)).

In this paper, we propose a theory involving heterogeneous firms facing costs to utilizing previously installed capital. The costs endogenously constrain the capacity each firm utilizes, and these constraints change with economic conditions. We investigate the impact of these capacity constraints on aggregate production and productivity when the economy is driven by aggregate and idiosyncratic demand shocks. We devise a model of demand shocks and endogenous capacity constraints that is qualitatively consistent with these observations. We then use the mentioned stylized facts in combination with data on firms' input utilization rates in order to learn about the relevance of demand shocks for aggregate fluctuations. To assess this, we estimate the model using micro data and standard Bayesian techniques. We then compare to which extent it replicates the described stylized facts. We find that for our preferred calibration, the model exhibits asymmetry on the order of what we observe in the data.

Capacity utilization, and its relationship to business cycles, has been the subject of a long strand of research in macroeconomics. Hansen and Prescott (2005), motivated by the observation that many firms typically have idle capacity, construct a model in which a representative firm chooses capacity one period ahead, and finds that business cycles in such a model are asymmetric. Gilchrist and Williams (2005) investigate the use of "putty-clay" technology by firms, in which firms freely choose their capital a period in advance, and can

only utilize that capital during production. They focus on the construction of a tractable aggregate production function, and find that business cycles are asymmetric as well. The model in this paper is closely related to this work.

This paper proceeds as follows. In the next section, we document our three stylized facts. Next, we propose our model, and discuss its important properties. In the next section, we estimate the key parameters of our model, simulate the model using these parameters, and discuss its quantitative performance. The last section concludes.

2 Three stylized facts relating to asymmetry of business cycles

Large deviations from trend in output are more likely to be negative The question whether business cycles are asymmetric is fairly old. As noted by McKay and Reis (2008), Mitchell (1927) characterized recessions as briefer and more violent than expansions. Starting with Neftci (1984) as well as DeLong and Summers, a large literature has investigated this question using more formal econometric techniques.

In general, results in the literature are mixed. Neftci (1984) finds evidence for asymmetry between increases and decreases of the US unemployment rate estimating an underlying Markov model. DeLong and Summers (1986) test skewness coefficients of output and employment data for the US and other industrialized countries. They do not find significant skewness in growth rates of output, nor significant differences in the length of expansions and contractions in output. Hamilton (1989) extends Neftci's work and estimates a two-state Markov process for output in which the recessionary and expansionary states have significantly different properties, a finding confirmed by later papers using similar methodology (Clements and Krolzig (2003), Hamilton (2005)). Sichel (1993) takes up DeLong and Summers (1986) nonparametric approach by looking at deepness of contractions and finds evidence of negative skew in levels of output and employment. Bai and Ng (2005) derive limiting distributions for coefficients of skewness and kurtosis under serial correlation and cannot reject the null of zero skew in US output growth rates. After nonparametrically estimated business cycle turning points, McKay and Reis (2008) compare average duration

and growth during expansions and recessions, and do not find significant evidence that recessions are shorter or more violent.

These findings suggest to us that one should be specific in defining which aspect of potential asymmetry one is looking at, a point also made by McKay and Reis (2008). Our general reading of the literature is as follows: Evidence for skewness in output growth rates, which seem to be considered in most papers, is largely absent. There is some evidence for skewness in output levels. (Along both dimensions, employment seems to exhibit stronger asymmetry than output.) Additionally, for hidden-state Markov models, non-linear parameterizations fit the data significantly better than linear ones.

Our focus is on the claim that large cyclical swings in output are more likely to be negative than positive — in other words, production will spend more time far below trend than far above. This is related to negative skewness of output levels, for which there is some evidence (Sichel (1993)).

We now report some additional observations about the largest absolute deviations of the cyclical component of output. To this purpose, we take a measure of US production, detrend it, and then look at the largest observations of the detrended series in absolute value. Specifically, for some integer N we count in how many of the N periods of largest absolute deviations output was above trend vs below trend. We also compare the mean of the $N/2$ most positive deviations to the $N/2$ most negative deviations. Finally, we also compute the coefficient of skewness for the detrended series as a whole (so that this statistic is independent of N). We repeat this process for a number of different specifications, where a specification consists of a measure of output, a time span, a trend filter, and an integer N .

We report results in table 1. Our baseline cases use HP-filtered postwar data and often constitutes the weakest case in terms of differences between expansions and recessions since the HP filter tends to attribute parts of the cyclical movement into the trend at the edges of the sample. For almost all specifications we find that large deviations from trend tend to be negative.

The following discussion of the other two stylized facts is shorter, mainly because the evidence is not as mixed.

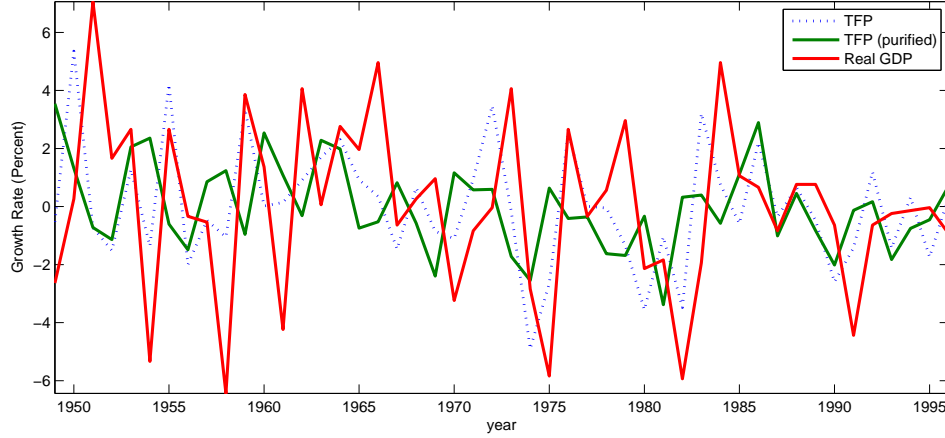
Table 1: Strong recessions more severe than strong expansions

Specification	# pos vs neg	mean pos vs neg	skewness
Quarterly GDP			
Baseline	16 vs 24	2.78% vs -3.68%	-0.60
$N = 20$	6 vs 14	3.12% vs -4.33%	-0.60
$N = 80$	40 vs 40	2.28% vs -2.87%	-0.60
until 2007	17 vs 23	2.28% vs -2.87%	-0.62
Linear filter	6 vs 34	7.99% vs -12.70%	-0.81
Rotemberg filter	6 vs 34	4.19% vs -5.68%	-0.33
Rotemberg filter, $N = 80$	29 vs 51	3.74% vs -5.14%	-0.33
Annual GDP			
Baseline	3 vs 7	3.20% vs -4.40%	-0.35
$N = 6$	0 vs 6	3.37% vs -4.83%	-0.35
$N = 20$	13 vs 7	2.99% vs -3.55%	-0.35
until 2007	4 vs 6	2.99% vs -3.55%	-0.35
from 1929	6 vs 4	16.69% vs -11.61%	+1.00
Linear filter	2 vs 8	7.29% vs -12.51%	-0.88
Linear filter from 1929	3 vs 7	20.50% vs -31.08%	-0.91
Rotemberg filter	1 vs 9	6.23% vs -13.50%	-0.87
Rotemberg filter from 1929	1 vs 9	16.15% vs -36.95%	-1.22
Monthly industrial production			
Baseline	50 vs 70	4.52% vs -5.90%	-0.65
$N = 40$	7 vs 33	5.48% vs -7.57%	-0.65
$N = 240$	124 vs 116	3.71% vs -4.45%	-0.65
until 2007	56 vs 64	4.39% vs -5.58%	-0.65
from 1919	54 vs 66	11.35% vs -13.59%	-0.55
Linear filter	33 vs 87	17.03% vs -22.69%	-0.52
Rotemberg filter	46 vs 74	7.47% vs -11.23%	-0.62

“# pos vs neg”: Out of the N periods with largest absolute value, how many were positive and how many were negative. “mean pos vs neg”: Mean of the $N/2$ largest periods vs mean of the $N/2$ smallest periods. “Skewness”: Coefficient of skewness defined as $E[(x - \mu)^3/\sigma^3]$.

For all three series in the baseline, N corresponds to a little less than 1/6 of observations, series were HP filtered and starting date is January 1949. “Quarterly GDP”: $N = 40$, end date 2014 : 4, HP(1600)-filtered. “Annual GDP”: $N = 10$, end date 2013, HP(400)-filtered. “Monthly industrial production”: $N = 120$, end date 2014/02, HP(10,000)-filtered. Alternative specifications differ from respective baseline only along listed dimensions.

Figure 1: GDP and TFP measures



Note: Real GDP series from FRED, along with TFP series and purified TFP series from Fernald (2012).

Simple TFP estimates are procyclical, but not so if corrected for utilization For this stylized fact we draw on Basu et al. (2006) and Basu et al. (2009) who discuss ways to improve the measurement of firm productivity. In particular, they construct a measure for aggregate technology that accounts for potential confounding influences of returns to scale, imperfect competition, aggregation across sectors and, especially relevant for us, utilization rates of factor inputs. Figure 1 displays annual output growth together with changes in a simple measure of TFP (the Solow Residual) as well as Basu, Fernald and Kimball’s ‘purified’ measure. The simple productivity measure is strongly procyclical: Correlation between output growth and simple TFP is 0.74. The improved technology measure does not exhibit this association with aggregate production; in fact purified TFP appears to be almost completely acyclical as its correlation with (contemporary) output growth is 0.02.

Since the mechanism we consider hinges strongly on the effect of adjustment in factor input utilization, we recalculate the mentioned coefficients of correlation using data provided by John Fernald¹ (see Fernald (2012)). This dataset provides among other things a TFP measure that *only* corrects for intensity of capital and labor utilization, allowing us to check if utilization is indeed relevant for the difference in cyclicity between the simple and

¹Data available at www.frbsf.org/economic-research/economists/jferald/quarterly_tfp.xls

the purified productivity measure (or if instead the difference stems mainly from the other ‘purifying’ steps taken by Basu, Fernald and Kimball). Additionally, it spans 15 more years at the end of the sample. Again, simple TFP is strongly procyclical with a correlation of 0.83 whereas utilization-corrected TFP is acyclical with a coefficient of -0.03 .

Dispersion of firm productivity increases in recessions Our third fact is connected to a range of findings in the literature that relate recessions to increased cross-sectional dispersion among firms along several dimensions. Eisfeldt and Rampini (2006) show that capital productivity is more dispersed in recessions. Bloom (2009) and Bloom et al. (2012) show in seminal papers that shocks to the variance of firm productivity can cause drops in output; they include empirical evidence relating dispersion in sales growth, innovations to plant profitability, and sectoral output to times of low aggregate production.

Directly related to levels of firm productivity, Kehrig (2011) finds that the distribution of plant revenue productivity becomes wider in recessions; Bachmann and Bayer (2011) reach a similar result for innovations to the Solow residual in a dataset of German firms.

3 Model

We construct a simple monopolistic competition model in which a final good is produced with a variety of intermediate inputs. Each of the intermediate input producers must choose the price of their good and their capital capacity one period ahead of when they decide to produce. At the time of production, these firms are free to choose their labor input and *utilized* capital input. A convex cost of capital utilization ensures that firms choose interior values for their utilization rates.

This is a partial equilibrium model, in which wages, rental rates and, importantly, demand functions are independent of the outcomes in the intermediate goods sector.² Importantly, this shuts down an interaction between binding constraints on intermediate goods firms and aggregate demand: An increasing share of constrained intermediate goods producers will potentially increase demand for other, non-constrained firms’ goods and will

²This setup could, for example, be rationalized as part of a general equilibrium in which there is a global final goods aggregator, a global household, and the intermediate goods sector represents a country’s firms that jointly have mass 0 in final goods production.

also lower final goods production. Both these effects can have feedback effects on the share of constrained firms and thus deprive the Dixit-Stiglitz solution of its usual simplicity. For this reason we leave generalization to general equilibrium for future work.

3.1 Final Goods Firms

There is a competitive sector of final goods producers which acts as a CES aggregator, that is, production of final goods Y_t requires purchasing a basket of intermediate inputs of differing varieties $\{y_{it}^d\}_i$ at prices $\{p_{it}\}_i$. Each final goods firm's production function is given by

$$Y_t = \left[\int (A_t b_{it})^{\frac{1}{\sigma}} y_{it}^{\frac{\sigma-1}{\sigma}} di \right]^{\frac{\sigma}{\sigma-1}} \quad (1)$$

where σ is the (constant) elasticity of substitution between inputs. A_t and b_{it} are random (AR(1)) aggregate and variety-specific final good technology shocks, respectively. We will alternatively refer to these as "demand shocks" throughout the remainder of the paper. Under the assumption that b_{it} is lognormally distributed, which we will maintain throughout the paper, the aggregate shock A_t just shifts the mean of b_{it} .

As will be made clear below, because of capital utilization costs, intermediate goods firms will not be willing to supply more than a (variety specific) quantity y_{it}^C . Also, the price of final goods is normalized to 1. Thus, in each period t , final goods firms solve

$$\max_{(y_{it})_i} Y_t - \int_i p_{it} y_{it} di$$

Define the expenditure for varieties $I_t = \int_i (A_t b_{it})^{\frac{1}{\sigma}} p_{it} y_{it} di$, and price index $P_t^{1-\sigma} = \int_i (A_t b_{it})^{\frac{1-\sigma}{\sigma}} p_{it}^{1-\sigma} di$. Then

$$y_{it}^d = A_t b_{it} \frac{I_t}{P_t} \left(\frac{P_t}{p_{it}} \right)^\sigma$$

3.2 Intermediate Goods Firms

Intermediate goods producers transform capital and labor supplied by households into differentiated products that can be used as inputs by the final goods sector. Each producer

monopolistically supplies a single variety.

We depart from the standard model in three ways. First, each intermediate goods firm chooses the price of its variety one period in advance. Second, also one period in advance, each firm invests some of its earnings (in units of final goods) as installed capital for production in the next period. Third, firms pay a cost this period for utilizing installed capital that is increasing in the fraction of installed capital used.

We will allow firms to utilize capital beyond installed capital this period, so that the utilization rate is bigger than 100%; firms are essentially able, albeit at great cost, to "overwork" the machines installed in the factory last period as much as the care to. For the cost functions and parameterizations we consider, the cost of utilizing capital keeps firms from overworking installed capital in any period.

Specifically, call installed capital by the producer of intermediate input i in period t k_{it} , and the actual capital used in period t \tilde{k}_{it} , so that $\tilde{k}_{it} \leq k_{it}$. Then the producer of intermediate input i has a production function given by

$$y_{it} = \tilde{k}_{it}^\alpha l_{it}^{1-\alpha} \quad (2)$$

Installed capital is rented at the end of period $t - 1$ at the interest rate R_t , but producers do not repay households until the end of period t , as in the standard case. Capital depreciates by the rate δ after it is used, and firms maintain ownership of the undepreciated capital. Capital utilization comes at a cost given by the function $c(\tilde{k}, k)$. Thus, each firm i 's profit in period t is given by

$$\pi_{it} = p_{it}y_{it} - w_t l_{it} - R_t k_{it} - c(\tilde{k}_{it}, k_{it}) + (1 - \delta)k_{it} \quad (3)$$

We will assume that the capital utilization cost function is quadratic in the utilization rate:

$$c(\tilde{k}, k) = \chi \left(\frac{\tilde{k}}{k} \right)^2 k \quad (4)$$

Firm i maximizes the sum of per period profits, but since the only buyer of variety i is

the final goods producer, the firm cannot supply any more than y_{it}^d . The firm's maximization problem is given by

$$\begin{aligned} \max_{(p_{it}, k_{it})_{t=0}^{\infty}} E_t \sum_{i=1}^{\infty} \Lambda_{t,t+i} \max_{(\tilde{k}_{it}, l_{it})} \pi_{it} \\ \text{s.t. } y_{it} \leq y_{it}^d \end{aligned}$$

We solve the firm's problem in three steps. First, find the firm's choice of (\tilde{k}_{it}, l_{it}) that minimizes its cost of supplying y_{it} given $(p_{it}, k_{it}, A_t, b_{it})$, where A_t and b_{it} are (random) aggregate and variety-specific demand shocks. Second, find the firm's choice of y_{it} that maximizes its profits given $(p_{it}, k_{it}, A_t, b_{it})$. Third, find the choices of (p_{it}, k_{it}) that maximize the firm's expected profits, where expectations are taken over the random shocks (A_t, b_{it}) .

3.2.1 Firm Costs and Factor Demands

The firm chooses how much labor and capital to use in period t to minimize its period t costs.³ The rental costs of installed capital are sunk at the time t . Thus, given $(p_{it}, k_{it}, w_t, R_t, b_{it}, A_t)$ and a desired level of output y , the firm chooses (\tilde{k}_{it}, l_{it}) to minimize

$$\begin{aligned} \min_{(\tilde{k}_{it}, l_{it})} w_t l_{it} + \chi \left(\frac{\tilde{k}_{it}^2}{k_{it}} \right) \\ \text{s.t. } \tilde{k}_{it}^{\alpha} l_{it}^{1-\alpha} \geq y \end{aligned}$$

The solution to this problem is given by

$$\tilde{k}_{it} = y_{it}^{\frac{1}{2-\alpha}} \left(\frac{w_t}{2\chi} \frac{\alpha}{1-\alpha} \right)^{\frac{1-\alpha}{2-\alpha}} k_{it}^{\frac{1-\alpha}{2-\alpha}}$$

$$l_{it} = y^{\frac{1}{1-\alpha}} \tilde{k}_{it}^{-\frac{\alpha}{1-\alpha}}$$

With these factor demands, we can also find the firm's cost function:

³Minimizing period t costs minimizes the firm's total costs, since there adjusting these inputs have no effect on costs in other periods.

$$C(y) = w_t^{\frac{2-2\alpha}{2-\alpha}} \chi^{\frac{\alpha}{2-\alpha}} k_{it}^{-\frac{\alpha}{2-\alpha}} \left(\frac{2-\alpha}{2-2\alpha} \right) \left(\frac{2-2\alpha}{\alpha} \right)^{\frac{\alpha}{2-\alpha}} y^{\frac{2}{2-\alpha}}$$

3.2.2 Firm Supply

Using the cost function obtained above, we can state the firm's profit maximization problem and obtain the optimal level of output y_{it} . As the only buyer of variety i is the final goods producer, the firm cannot supply any more than y_{it}^d . The firm chooses y_{it} to maximize⁴

$$\begin{aligned} \max_{y_{it}} \quad & p_{it} y_{it} - C(y_{it}) \\ \text{s.t.} \quad & y_{it} \leq y_{it}^d \end{aligned}$$

We can now write the solution to the above profit maximization problem. Given $(p_{it}, k_{it}, b_{it}, A_t)$,

$$y_{it} = \begin{cases} y_{it}^d & \text{if } A_t b_{it} \leq \bar{b}_{it} \\ y_{it}^C & \text{else} \end{cases}$$

where

$$\begin{aligned} y_{it}^C &= \frac{\alpha}{2\chi(1-\alpha)} (1-\alpha)^{\frac{2-\alpha}{\alpha}} p_{it}^{\frac{2-\alpha}{\alpha}} w_t^{-\frac{2-2\alpha}{\alpha}} k_{it} \\ \bar{b}_{it} &= \left(\frac{P_t^{1-\sigma} p_{it}^\sigma}{I_t} \right) y_{it}^C \end{aligned}$$

The cutoff \bar{b}_{it} gives the lowest level of demand shock that ensures that the constraint $y_{it} \leq y_{it}^d$ does not bind. For all $A_t b_{it} > \bar{b}_{it}$, the marginal cost of supplying y_{it}^d exceeds the variety-specific price p_{it} , and therefore the firm would earn higher profits by only supplying y_{it}^C . This level of output acts as an endogenously determined limit on output (and therefore capital utilization); though this limit can differ across firms in the general model, in the solution we consider, it will be the same for all firms.

⁴The firm's profit function should also include the cost of repaying capital to the household, rented last period. Because this cost is sunk in period t , however, we ignore that cost in solving the period t output decision.

3.2.3 Firm Expected Profit

With the firm's supply function in hand, we can write down the state-by-state profit function, along with the expected profits the firm will earn given a choice of (p_{it}, k_{it}) and information about previous period random shocks $(A_{t-1}, (b_{i,t-1})_i)$.

After uncertainty is realized, profits (less capital rental costs) are given by

$$\pi_{it}^n(p_{it}, k_{it}, b_{it}, A_t) = \begin{cases} p_{it}y_{it}^d - C(y_{it}^d) & \text{if } A_t b_{it} \leq \bar{b}_{it} \\ p_{it}y_{it}^C - C(y_{it}^C) & \text{else} \end{cases}$$

Thus, the firm chooses p_{it} and k_{it} to solve

$$\max_{p_{it}, k_{it}} E_{t-1} \pi_{it}(p_{it}, k_{it})$$

where the profit agents expect to receive in period t , looking from the end of period $t - 1$, is

$$E_{t-1} \pi_{it}(p_{it}, k_{it}) = \int \left(\int_0^{\frac{\bar{b}_{it}}{A_t}} [p_{it}y_{it}^d - C(y_{it}^d)] p_b(x) dx + \int_{\frac{\bar{b}_{it}}{A_t}}^\infty [p_{it}y_{it}^C - C(y_{it}^C)] p_b(x) dx \right) dF_A$$

3.3 Solution

We begin by putting structure on the random shocks b_{it} and A_t . The distribution of the random shock A_t follows an AR(1) process with normal error, while the distribution of the idiosyncratic shock b_{it} follows an AR(1) process with lognormal error:

$$\begin{aligned} b_{it} &= \rho_b b_{i,t-1} + \epsilon_{bt} \\ \ln(A_t) &= \rho_A \ln(A_{t-1}) + \epsilon_{At} \\ \epsilon_{bt} &\sim_d \ln \mathcal{N}(\mu_b, \sigma_b) \\ \epsilon_{At} &\sim_d \mathcal{N}(\mu_A, \sigma_A) \end{aligned}$$

In the estimation and simulation that follow, we make two assumptions on the structure of the demand shocks A_t and b_{it} :

- $\rho_b = 0$ (aggregate errors are normal, idiosyncratic errors are lognormal)
- $\mu_b = -\frac{\sigma_b^2}{2}$ (normalize mean of shock to 1)

We solve the full set of equilibrium equations under the above assumptions in the appendix. Importantly, the iid assumption implies that the one-period ahead prices and installed capital choices p_{it} , k_{it} , and \bar{b}_{it} are identical across firms in every period. (We'll interchangeably call these common choices p_t and k_t .) In turn, this implies that the endogenous capacity constraint, y_{it}^C , and the corresponding demand shock cutoff \bar{b}_{it} , are identical for all firms.

Now consider the effect of an increase in the aggregate shock A_t . First, the cutoff \bar{b}_t doesn't depend on the aggregate shock A_t . This implies that when we increase A_t , essentially increasing the mean demand shock faced by any firm, we necessarily make more firms capacity constrained.

We can now show that the model qualitatively matches the second two stylized facts described above, though we do not lay out the proofs in detail here. The second fact states that while measured TFP is procyclical, "purified" TFP (controlling for factor utilization) is less procyclical. The analog for both of these quantities is

$$\begin{aligned} \text{Measured TFP: } \frac{y_{it}}{k_{it}^\alpha l_{it}^{1-\alpha}} &= \left(\frac{\tilde{k}_{it}}{k_{it}} \right)^\alpha \\ \text{Purified TFP: } \frac{y_{it}}{\tilde{k}_{it}^\alpha l_{it}^{1-\alpha}} &= 1 \end{aligned}$$

The first quantity is just a function of the capital utilization rate; using the aggregate (mean) capital utilization rate from the appendix, it can be shown that the derivative of the above function with respect to A_t is positive, and therefore increasing in A_t . Since we use A_t to stand-in as the underlying source of business cycle fluctuations, we can say that this quantity is procyclical. The second quantity, on the other hand, is totally acyclical.

The third fact states that productivity dispersion increases in recessions. One way to show this is to show that the coefficient of variation of measured TFP is an decreasing function of A_t , where

$$CV_{TFP} = \frac{\sigma_{TFP}}{\mu_{TFP}} = \frac{\int \left(\frac{\tilde{k}_{it}}{k_{it}} \right)^{2\alpha} di}{\left(\int \left(\frac{\tilde{k}_{it}}{k_{it}} \right)^{\alpha} di \right)^2} - 1$$

Again, though we do not show it here, the derivative of the above function with respect to A_t is negative, and therefore decreasing in A_t .

4 Estimation

In the previous section, we showed that the model qualitatively implies a connection between capacity utilization and business cycle asymmetry. We now want to investigate to which extent the capacity and capital utilization that we observe in the data is quantitatively relevant, in that it implicitly generates a significant portion of observed business cycle asymmetry. We do this by first estimating model parameters only with the capacity and capital utilization data, and then simulating the calibrated model. In this section, we outline the data we use for the calibration, then outline our calibration procedure and results.

4.1 Data

We use two aggregate time series on the manufacturing sector for the estimation: a measure of capacity utilization and another for capital utilization. All data are annual, and we restrict our attention to the range 1974-2004 for which the capital utilization data is available. We obtain aggregate capacity utilization across surveyed firms from FRED⁵ from the Quarterly Survey of Plant Capacity Utilization conducted by US Census Bureau (2014). The data surveys plant managers across several sectors of the economy. Managers are asked to estimate the current market value of their production as a percentage of the value of production that would occur if plants were operating at "full capacity", i.e. using all machinery

⁵Available at <https://research.stlouisfed.org/fred2/series/MCUMFN>

in place, and as much labor/fuel/etc. as necessary to operate it.

We obtain capital utilization data from Gorodnichenko and Shapiro (2011)⁶, using their preferred measure of plant hours per week. This measure uses information on labor hours to estimate how intensively a plant’s machines are being used.

For the output measure we use data on real GDP from FRED⁷. We detrend the annual production series using an HP(100) filter. Since plant hours per week and capacity utilization are naturally bounded from both sides, it is not clear that they should be detrended. We choose to HP-filter them with a high penalty parameter of 10,000, in order to remove very long-run trends. Finally, for output and workweek of capital, we use log deviations from trend since the levels of these variables are not informative in the context of our model. However, we continue to express capacity utilization in levels, since the level of its long-run average of around 78% corresponds naturally to its equivalent in the model.

4.2 Estimation Procedure

We estimate four parameters: the cost parameter χ , the variance of the idiosyncratic demand process σ_b , as well as persistence and variance of the aggregate demand shock ρ_A and σ_A , respectively.

Given a set of parameters (χ, σ_b) , and the mean wage w_t , we use the model to calculate mean capital utilization and capacity utilization. We then match this against the capital utilization and capacity utilization from the data.

For capacity utilization we refer to full capacity as y_{it}^C , the maximum output an individual firm is willing to supply when demand is very high. While a more natural counterpart to the survey question would be to compute output at some level of maximum capital utilization, it is not entirely clear what this would mean in the context of the model.

We obtain our estimates using a standard Bayesian approach; the procedure (see Fernández-Villaverde (2010)) is described in more detail in the appendix. The results are displayed in table 2. The first three parameters are chosen as standard in the literature. For the elasticity of substitution σ , the estimates in the literature range from around 4 to 10. For now we choose a value on the higher end of estimates (giving firms less monopolistic power),

⁶ Available at <http://www-personal.umich.edu/~shapiro/data/SPC/>

⁷ Available at <https://research.stlouisfed.org/fred2/series/IPMAN>

Table 2: Parameters

Parameter		Value
<i>Set outside estimation</i>		
Effective Capital Share	α	0.33
Discount Factor	β	0.96
Depreciation Rate	δ	0.1
Elasticity of Substitution	σ	10
Autocorr. Demand Shock	ρ_b	0
<i>Estimates</i>		
Utilization Cost	χ	0.299
Std. Dev. Demand Shock	σ_b	4.67
Autocorr. Agg Shock	ρ_A	0.496
Std. Dev. Agg Shock	σ_A	0.041

and note that it would be a good idea to include the parameter in future estimations. The last parameter set outside the estimation is ρ_b , the persistence of the idiosyncratic demand shock. In the context of the current model, this is without loss of generality: since firm decisions are linear in expected demand, the distribution of $E[b_i]$ among firms does not matter (nor are we using data identifying this distribution). The next four parameters represent the median values of their respective marginal posterior distribution; we give further information about the posteriors in the appendix. While the variance of the idiosyncratic demand shock may seem high, it is worth keeping in mind that the mean of the distribution over b_i is kept normalized to 1 (the lognormal distribution has parameters $(-\sigma_b^2/2, \sigma_b)$). This means that for most firms the realization of their demand shock actually becomes smaller as σ_b increases.

4.3 Results

We now assess to which extent the model reproduces the three stylized facts under the calibrated parameters found in the previous section.

To do so, we use the the median estimates for χ and σ_b and compare the effects of symmetric aggregate demand shocks on output, aggregate (non-corrected) TFP, and dispersion in average costs of production.

Table 3: Quantitative Performance

<i>Business Cycle Asymmetry</i>		
% Δ from SS level	20% change in A	-20% change in A
ΔKU	0.87	-0.92
ΔY	1.12	-1.19
	Increase	Decrease
Shock needed for 1% Change in Y	18	-17
Shock needed for 5% Change in Y	93	-88
		Skewness in Y
Model		-0.37
Data: Rotemberg Filter		-0.33
<i>Productivity Dispersion</i>		
Dispersion	Correlation with Y	
Measured TFP Variance (model)	-0.98	
TFPR Variance (data)	-0.5	

Note: Performance of the model with respect to business cycle asymmetry and productivity dispersion. First, we list percent changes from steady state for production-weighted aggregate capital utilization (KU), capacity utilization (YU), and output (Y). Next, we list the shock to the aggregate shock A that would be needed to produce a 1% increase and decrease in Y , respectively, and then again for a 10% change. Next, we list the skewness of a simulated output series given the standard deviation for aggregate shock process A that generates the observed standard deviation of output. Last, we list correlation of productivity dispersion with output, both for a measured TFP in the model and estimates from Kehrig (2011).

The quantitative comparison results are given in table 3. We only report demand-weighted aggregates, to reflect final output as the product of the final goods aggregator. We first consider the impact of a change in the aggregate demand shock A , which corresponds to a change in the mean of the idiosyncratic demand shocks b_i . When A increases, more firms have higher demand shocks, but by the same token, more firms are producing at the endogenous output limit y_t^C .

We first shock the model with a shock equal to the calibrated standard deviation of $\sigma_{AY} = 0.2$. Generally, even for a large change in A , the model delivers small changes in all aggregate variables. The first three rows of the table show this: a 20% shock to aggregate

demand relatively modest changes in output, capital utilization, and capacity utilization. We believe these magnitudes have something to do with the large estimated σ_b - since the distribution of demand shocks is already wide with $A = 1$, a small change in A doesn't affect the demand of many firms.

That said, when the model is calibrated to match the magnitude of observed changes in output, it does exhibit asymmetry in output. Given the same percentage magnitude aggregate shock, output decreases (as a percentage of its steady state value) by more than it increases (1.19% decrease vs 1.12% increase). This difference is small, but as the next rows show, it does have significant implications for the frequency of large busts versus large booms.

Under the calibration for the aggregate shock process A_t , we know the distribution of possible aggregate shocks. We can use the calibrated parameters to simulate the model for several demand shocks and calculate the implied skewness in aggregate output in the model. We see skewness of output in the model is on the order of what we see in the data; for example, the skewness in the filtered quarterly GDP data is a little less than the skewness implied by the model.

The model also seems to predict more skewness at higher levels of volatility. In the second two rows, we list the level of A_t necessary to induce a 1% increase and decrease in GDP, and then a 5% increase and decrease in GDP. At the 1% level, the necessary aggregate shocks are very close together in magnitude, only differing by a percentage point. At the 5% level, the difference between the negative shocks increases to 5 percentage points. We take this to imply that the utilization effects we capture in the model may be more pronounced for large changes in output.

Finally, the model also predicts significantly higher TFP dispersion in busts. Using the formula

$$\widetilde{TFP}_{it} = \frac{y_{it}}{k_{it}^\alpha l_{it}^{1-\alpha}} \quad (5)$$

We calculate the variance of \widetilde{TFP}_{it} for many draws of the aggregate shock with $\sigma_A = 0.2$. We then calculate the correlation between the variance of TFP and output; we find it to

almost perfectly negatively correlated. The corresponding estimate from Kehrig (2011) is much (unsurprisingly) significantly smaller, on the order of 0.5.

As shown above, the analog of "purified" TFP will always be acyclical in this model, while simple TFP will be procyclical, so we cannot meaningfully add to our model's performance on this measure in this quantitative exercise. We can speak to the model's general performance, however, by comparing the implied correlation between output growth and simple TFP.

5 Conclusion

We present a model of demand shocks which can qualitatively explain three business cycle observations: Deep recessions, procyclicality of the Solow residual while purified TFP is acyclical, and countercyclical dispersion in firms' Solow residuals and average costs. The main assumption is that firms set prices before the level of their demand is realized. We then demonstrate how to estimate the model using time-series data on the level of capacity utilization as well as the comovement of output, capacity utilization and capital utilization. Further research is needed to derive the general equilibrium solution to the model as well as adapting the model to allow for a more informed estimation using more data.

References

- Bachmann, R. and C. Bayer (2011). Investment dispersion and the business cycle. *NBER Working Paper* (w16861).
- Bai, J. and S. Ng (2005). Tests for Skewness, Kurtosis, and Normality for Time Series Data. *Journal of Business & Economic Statistics* 23(1), 49–60.
- Basu, S., J. Fernald, J. Fisher, and M. Kimball (2009). Sector-Specific Technical Change.
- Basu, S., J. G. Fernald, and M. S. Kimball (2006). Are Technology Improvements Contractionary? *American Economic Review* 96(5), 1418–1448.
- Bloom, N. (2009). The Impact of Uncertainty Shocks. *Econometrica* 77(3), 623–685.

- Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta Eksten, and S. Terry (2012). Really Uncertain Business Cycles. *NBER Working Paper* (w18245).
- Clements, M. and H. Krolzig (2003). Business Cycle Asymmetries. *Business & Economic Statistics* 21(1), 196–211.
- DeLong, J. and L. Summers (1986). Are Business Cycles Symmetric? *NBER Working Paper Series* 1444.
- Eisfeldt, A. L. and A. a. Rampini (2006). Capital Reallocation and Liquidity. *Journal of Monetary Economics* 53(3), 369–399.
- Fernald, J. (2012). A Quarterly, Utilization-Adjusted Series on Total Factor Productivity. *Manuscript, Federal Reserve Bank of San Francisco*.
- Fernández-Villaverde, J. (2010). The econometrics of DSGE models. *SERIEs* 1(1-2), 3–49.
- Gilchrist, S. and J. C. Williams (2005). Investment, capacity, and uncertainty: a putty-clay approach. *Review of Economic Dynamics* 8(1), 1–27.
- Gorodnichenko, Y. and M. Shapiro (2011). Using the Survey of Plant Capacity to Measure Capital Utilization. *US Census Bureau CES*, 11–19.
- Hamilton, J. (1989). A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica* 57(2), 357–384.
- Hamilton, J. D. (2005). What’s Real about the Business Cycle? Technical Report w11161, National Bureau of Economic Research.
- Hansen, G. D. and E. C. Prescott (2005). Capacity constraints, asymmetries, and the business cycle. *Review of Economic Dynamics* 8(4), 850–865.
- Kehrig, M. (2011). The Cyclicalilty of Productivity Dispersion. *US Census Bureau Center for Economic Studies Paper No. CES-WP-11-15*.
- McKay, A. and R. Reis (2008). The Brevity and Violence of Contractions and Expansions. *Journal of Monetary Economics* 55(4), 738–751.

- Mitchell, W. C. (1927). The Processes Involved in Business Cycles. In *Business Cycles: the Problem and its Setting*, pp. 1–60. NBER.
- Neftci, S. N. (1984). Are Economic Time Series Asymmetric over the Business Cycle? *Journal of Political Economy*, 307–328.
- Sichel, D. E. (1993). Business Cycle Asymmetry: a Deeper Look. *Economic Inquiry* 31(2), 224–236.
- US Census Bureau (2014). Quarterly survey of plant capacity utilization. Available online at `bhs.econ.census.gov/bhs/pcu/index.html`.

A Solution

In this section, we list the full set of equilibrium equations under the assumptions that idiosyncratic demand shocks b_{it} are iid, and the shock process has mean 1, so that

$$b_{it} = \epsilon_{bt}$$

$$\epsilon_{bt} \underset{d}{\sim} \ln \mathcal{N}\left(-\frac{\sigma_b^2}{2}, \sigma_b\right)$$

The first assumption implies that the one-period ahead prices and installed capital choices p_{it} and k_{it} are identical across firms in every period. We’ll interchangeably call these common choices p_t and k_t . In turn, this implies that several quantities in the model also become identical across firms, e.g. $\bar{b}_{it} = \bar{b}_t \forall i$. We present only the simplified equations below.

Since this is a partial equilibrium model, both wages w_t and rental rates for capital R_t are treated as given.

Final Goods Firms

$$Y_t = \left(\left[\int_0^{\bar{b}_t} (A_t b_{it})^{\frac{1}{\sigma}} y_{it}^d \frac{\sigma-1}{\sigma} di \right] \right)^{\frac{\sigma}{\sigma-1}} \quad (6)$$

$$Y_t = I_t \quad (7)$$

$$P_t = \left(\int_0^\infty (A_t b_{it})^{\frac{1-\sigma}{\sigma}} p_{it}^{1-\sigma} di \right)^{\frac{1}{1-\sigma}} = A_t^{\frac{1}{\sigma}} p_t \quad (8)$$

$$I_t = \int_0^\infty (A_t b_{it})^{\frac{1}{\sigma}} p_{it} y_{it}^d di = p_t \int_0^\infty (A_t b_{it})^{\frac{1}{\sigma}} y_{it}^d di \quad (9)$$

$$y_{it}^d = \frac{I_t}{P_t} \left(\frac{P_t}{p_{it}} \right)^\sigma A_t b_{it} = \frac{I_t}{P_t} A_t^2 b_{it} \quad (10)$$

Intermediate Goods Firms

$$y_{it}^C = y_t^C = C_1 p_t^{\frac{2-\alpha}{\alpha}} w_t^{-\frac{2-2\alpha}{\alpha}} k_t \quad (11)$$

$$\bar{b}_{it} = \bar{b}_t = \frac{P_t}{I_t} y_t^C \quad (12)$$

$$\tilde{k}_{it} = C_0^{\frac{1-\alpha}{2-\alpha}} k_{t-1}^{\frac{1-\alpha}{2-\alpha}} w_t^{\frac{1-\alpha}{2-\alpha}} y_{it}^{\frac{1}{2-\alpha}} \quad (13)$$

$$l_{it} = y_{it}^{\frac{1}{1-\alpha}} \tilde{k}_{it}^{-\frac{\alpha}{1-\alpha}} \quad (14)$$

$$C(y) = C_2 \frac{w_t^{\frac{2-2\alpha}{2-\alpha}}}{k_t^{\frac{\alpha}{2-\alpha}}} y^{\frac{2}{2-\alpha}} \quad (15)$$

$$C_0 = \frac{\alpha}{2\chi(1-\alpha)} \quad (16)$$

$$C_1 = \frac{\alpha}{2\chi(1-\alpha)} (1-\alpha)^{\frac{2-\alpha}{\alpha}} \quad (17)$$

$$C_2 = \left(\frac{2-\alpha}{2-2\alpha} \right) \left(\frac{\alpha}{2\chi(1-\alpha)} \right)^{\frac{-\alpha}{2-\alpha}} \quad (18)$$

$$C_3 = \frac{2C_1^{\frac{\alpha}{2-\alpha}} - 2C_2}{\alpha C_1^{\frac{\alpha}{2-\alpha}}} \quad (19)$$

$$h_{1t}(p_{it}, k_{it}) = \int \left(\int_0^{\bar{b}_t} [p_{it} y_{it}^d - C(y_{it}^d)] p_b(x) dx \right) dF_A \quad (20)$$

$$h_{2t}(p_{it}, k_{it}) = \int \left(\int_{\bar{b}_t}^\infty [p_{it} y_{it}^C - C(y_{it}^C)] p_b(x) dx \right) dF_A \quad (21)$$

$$h_t = E_{t-1} \pi_{it}(p_{it}, k_{it}) = h_{1t} + h_{2t} \quad (22)$$

The FOC for p_{it} is given by

$$\frac{\partial h_t}{\partial p_{it}} = \frac{\partial h_{1t}}{\partial p_{it}} + \frac{\partial h_{2t}}{\partial p_{it}} = 0$$

And for k_t

$$\frac{\partial h_t}{\partial k_{it}} = \frac{\partial h_{1t}}{\partial k_{it}} + \frac{\partial h_{2t}}{\partial k_{it}} = R_t - (1 - \delta)$$

Exogenous Shock Processes

$$b_{it} \sim \ln \mathcal{N}\left(-\frac{\sigma_b^2}{2}, \sigma_b\right) \quad (23)$$

$$\ln(A_t) = \rho_A \ln(A_{t-1}) + \epsilon_{At} \quad (24)$$

$$\epsilon_{At} \sim_d \mathcal{N}(0, \sigma_A) \quad (25)$$

B Estimation Details

B.1 Moment Conditions

The aggregate (mean) capital utilization rate can be found with

$$\int \frac{\tilde{k}_{it}}{k_t} di = \int \frac{(C_0 w_t)^{\frac{1-\alpha}{2-\alpha}}}{k_t^{\frac{1-\alpha}{2-\alpha}}} y_{it}^{\frac{1}{2-\alpha}} di \quad (26)$$

$$= \frac{(C_0 w_t)^{\frac{1-\alpha}{2-\alpha}}}{k_t^{\frac{1-\alpha}{2-\alpha}}} \left(\frac{A_t I_t}{P_t} \right)^{\frac{1}{2-\alpha}} \left(\int_0^{\frac{\bar{b}_t}{A_t}} x^{\frac{1}{2-\alpha}} p_b(x) dx \right) + C_0^{\frac{1-\alpha}{2-\alpha}} C_1^{\frac{1}{2-\alpha}} w_t^{-\frac{1-\alpha}{\alpha}} p_t^{\frac{1}{\alpha}} \left(\int_{\frac{\bar{b}_t}{A_t}}^{\infty} p_b(x) dx \right) \quad (27)$$

We can express the aggregate (mean) capacity utilization rate by

$$\int \frac{y_{it}}{y_{it}^C} di = \left(\frac{A_t^2 I_t}{A_t^{\frac{1}{\sigma}} p_t^{\frac{2}{\alpha}} k_t} \right) w_t^{\frac{2-2\alpha}{\alpha}} \left(\int_0^{\frac{\bar{b}_t}{A_t}} x p_b(x) dx \right) + 1 \quad (28)$$

In the estimation, we assume that wages w_t are given. Then as long as p_{it} and k_{it} are only functions of the parameters (χ, σ_b) and the aggregate expenditure $I_t = Y_t$. The aggregate expenditure is itself only a function of A_t . Then the above capital utilization and

capacity utilization rates are also only functions of (χ, σ_b) and, through its effect on I_t , also a function of A_t . We can write p_{it} and k_{it} as functions of (χ, σ_b) using the FOCs from every firm's (common) one period ahead decisions.

To get the FOCs for the firm's one period ahead decisions, first take the partial derivatives of the components of the expected profit function:

$$\begin{aligned}\frac{\partial h_{1t}}{\partial p_{it}} &= \left(\int_0^{\bar{b}_t} \left(y_{it}^d + (p_{it} - C'(y_{it}^d)) \left(\frac{\partial y_{it}^d}{\partial p_{it}} \right) \right) p_b(x) dx \right) + [p_{it} \bar{y}_{it}^d - C(\bar{y}_{it}^d)] p_b(\bar{b}_t) \frac{\partial \bar{b}_t}{\partial p_{it}} \\ \frac{\partial h_{2t}}{\partial p_{it}} &= \left(y_{it}^C + (p_{it} - C'(y_{it}^C)) \left(\frac{\partial y_{it}^C}{\partial p_{it}} \right) \right) \left(\int_{\bar{b}_t}^{\infty} p_b(x) dx \right) - [p_{it} y_{it}^C - C(y_{it}^C)] p_b(\bar{b}_t) \frac{\partial \bar{b}_t}{\partial p_{it}}\end{aligned}$$

$$\begin{aligned}\frac{\partial h_{1t}}{\partial k_{it}} &= \left(\frac{\alpha}{(2-\alpha)k_{it}} \int_0^{\bar{b}_t} C(y_{it}^d) p_b(x) dx \right) + [p_t \bar{y}_{it}^d - C(\bar{y}_{it}^d)] p_b(\bar{b}_t) \frac{\partial \bar{b}_t}{\partial k_t} \\ \frac{\partial h_{2t}}{\partial k_{it}} &= \left((p_t - C'(y_t^C)) \frac{\partial y_t^C}{\partial k_t} + \frac{\alpha}{(2-\alpha)k_{it}} C(y_t^C) \right) \left(\int_{\bar{b}_t}^{\infty} p_b(x) dx \right) - [p_t y_t^C - C(y_t^C)] p_b(\bar{b}_t) \frac{\partial \bar{b}_t}{\partial k_t}\end{aligned}$$

where $\bar{y}_{it}^d = y_{it}^d(\bar{b}_t)$.

Expanding further,

$$\frac{\partial h_t}{\partial p_{it}} = (1-\sigma) \frac{I_t}{P_t} + \sigma (y_t^C)^{\frac{-\alpha}{2-\alpha}} \left(\frac{I_t A_t^2}{P_t} \right)^{\frac{2}{2-\alpha}} \left(\int_0^{\frac{\bar{b}_t}{A_t}} x^{\frac{2}{2-\alpha}} p_b(x) dx \right) + y_t^C \left(\int_{\frac{\bar{b}_t}{A_t}}^{\infty} p_b(x) dx \right) \quad (29)$$

$$\frac{\partial h_t}{\partial k_{it}} = \frac{\alpha}{2-\alpha} C_2 \left(\frac{I_t A_t^2}{P_t} \frac{w_t^{1-\alpha}}{k_{it}} \right)^{\frac{2}{2-\alpha}} \left(\int_0^{\frac{\bar{b}_t}{A_t}} x^{\frac{2}{2-\alpha}} p_b(x) dx \right) + \left(\frac{\alpha y_t^C}{p_{t-1} k_{t-1}} \right) \left(\int_{\frac{\bar{b}_t}{A_t}}^{\infty} p_b(x) dx \right) \quad (30)$$

We can now use equations (27) and (28) together with (29) and (30) to obtain moment conditions.

B.2 Estimation Procedure

Call the two model moments $f_{1t}(\chi, \sigma_b | A_t) = \int \frac{\tilde{k}_{it}}{\tilde{k}_t} di$, $f_{2t}(\chi, \sigma_b | A_t) = \int \frac{y_{it}}{y_{it}^C} di$, and their observed analogs m_{1t} and m_{2t} . Through their dependence on output, these functions also depend on the aggregate shocks A_t . Roughly, then, the estimation procedure proceeds in two steps: first, given a set of possible parameters $(\hat{\chi}, \hat{\sigma}_b)$, calculate the model moments for the entire distribution of possible shocks A_t . Second, calculate the expected difference between the model and corresponding data moments, given the prior distribution of shocks A_t ; if the difference is large, adjust the prior distribution, and repeat until the difference is minimized. Third, in an outer step, repeat this calculation for all periods t . Fourth, calculate a likelihood of observing the entire sequence of data moments, given your minimized shock distribution in each period. Last, repeat the entire process for other possible parameter sets $(\hat{\chi}, \hat{\sigma}_b)$ until the likelihood is minimized.

In solving the inner problem of adjusting the probability distribution of shocks A_t , we make use of the fact that the A_t has a parametric prior distribution with mean 0, and so reduce the problem of finding the best distribution to one of finding the best σ_A . In order to force the algorithm to search over many parameter sets, we assume that the observations m_{1t} and m_{2t} are observed with measurement error. We employ a Bayesian particle filter to solve the inner problem and generate a likelihood function. We then minimize the likelihood function using the Metropolis-Hastings algorithm; both of these are described in more detail in Fernandez-Villaverde (2009).

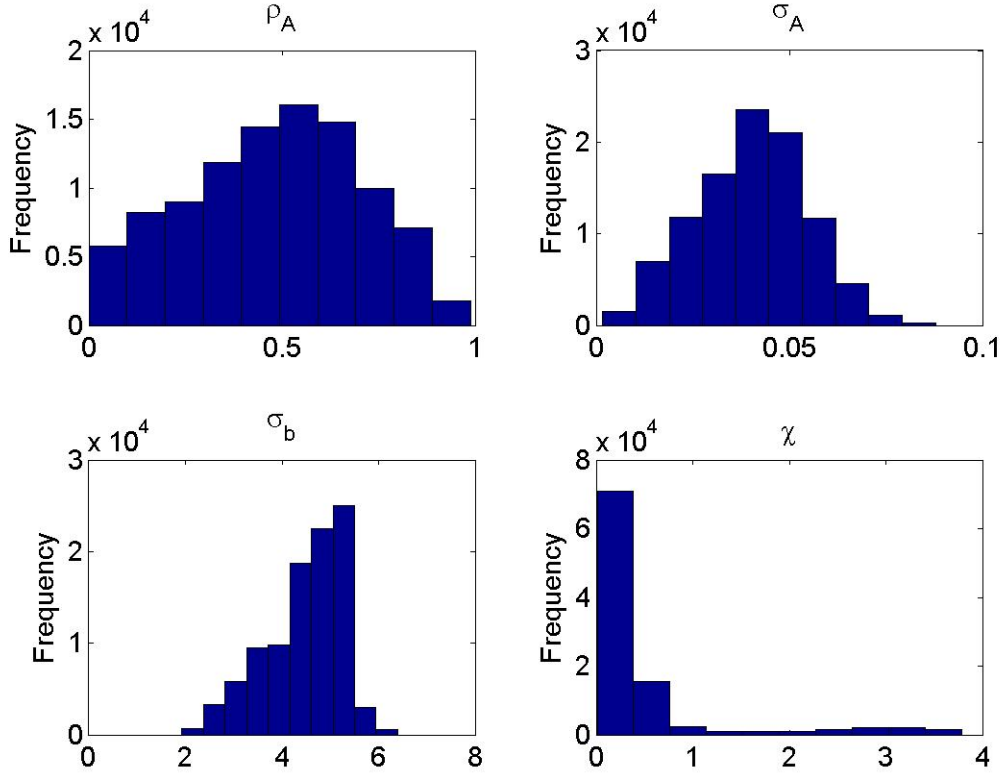
B.3 Results

We assume the following prior distributions on each of the estimated parameters:

Table 4: Assumed Priors

Parameter		Prior
Utilization Cost	χ	$\ln\mathcal{N}(0, 1)$
Std. Dev. Demand Shock	σ_b	$\ln\mathcal{N}(-2, 1)$
Autocorr. Agg Shock	ρ_A	$Uniform(0, 1)$
Std. Dev. Agg Shock	σ_A	$\ln\mathcal{N}(-2, 1)$

Figure 2: Posterior Distributions: Parameter Estimates



We list our parameter estimates in the main text. Here, we plot the approximate posterior distribution for each of the four parameters of interest.